

WHAT IS CLAIMED IS:

1. An interconnect apparatus, comprising:
a plurality of nodes; and
a plurality of interconnect lines selectively coupling
the nodes in a multiple level structure, the
multiple level structure being arranged to include:
a plurality of $J+1$ levels in a hierarchy of levels
T arranged from a level T equal to 0 to a
level T equal to J;
a plurality of 2^{J-T} rings in each level T; and
a plurality of 2^TK nodes in a ring.
2. An apparatus according to Claim 1 wherein a node A
on a level T greater than 0 and less than J has a plurality
of interconnections including:
an input interconnection from a node B on the level T;
an input interconnection from a node C on a level T+1;
an output interconnection to a node D on the level T;
and
an output interconnection to a node E on a level T-1.
3. An apparatus according to Claim 2 wherein a node A
on a level T greater than 0 and less than J has a plurality
of interconnections including:
a control input interconnection from the node F on the
level T-1; and
a control output interconnection to the node G on the
level T+1.
4. An apparatus according to Claim 2 wherein a node A
on a level T greater than zero and less than J has a
plurality of interconnections further including:
an input interconnection from a node H on a level T-2;
and
an output interconnection to a node I on a level T+2.

5. An apparatus according to Claim 4 wherein a node A on a level T greater than zero and less than J has a plurality of interconnections further including:

- a control input interconnection from a node J on a level T+2; and
- a control output interconnection to a node K on a level T-2.

6. An apparatus according to Claim 2 wherein at most one input interconnection of input connections B and C is active at one time.

7. An apparatus according to Claim 2 wherein at most one output interconnection of output connections D and E is active at one time.

8. An apparatus according to Claim 2 wherein messages communicated on the input interconnection from the node B on the level T have a higher priority than messages communicated on the input interconnection from the node C on the level T+1.

9. An apparatus according to Claim 2 wherein:
- a series of $2^T K$ sequential node A to node D interconnections on the level T traverses each of $2^T K$ nodes on one ring once.

10. An apparatus according to Claim 1 wherein the multiple level structure has a three-dimensional cylindrical topology in which each node has a location designated in three-dimensional cylindrical coordinates (r, θ , z) where radius r is an integer which specifies the cylinder number from 0 to J, θ is an integer which specifies the $2\pi/K$ spacing of nodes around the circular cross-section of a cylinder from 0 to K-1, and height z is a binary integer which specifies distance along the z-axis from 0 to $2^J - 1$.

11. An apparatus according to Claim 10 wherein:
a node $A(r, \hat{e}, z)$ is interconnected with an immediate predecessor node $B(r, (\hat{e}-1) \bmod K, H_r(z))$ on level r for receiving message data;
node $A(r, \hat{e}, z)$ is interconnected with a predecessor node $C(r+1, (\hat{e}-1) \bmod K, z)$ on level $r+1$ for receiving message data;
node $A(r, \hat{e}, z)$ is interconnected with an immediate successor node $D(r, (\hat{e}+1) \bmod K, h_r(z))$ on level r for sending message data;
node $A(r, \hat{e}, z)$ is interconnected with a successor node $E(r-1, (\hat{e}+1) \bmod K, z)$ on level $r-1$ for sending message data;
node $A(r, \hat{e}, z)$ is interconnected with a node $F(r-1, \hat{e}, H_r(z))$ on level $r-1$ for receiving a control input signal; and
node $A(r, \hat{e}, z)$ is interconnected with a node $G(r+1, \hat{e}, h_{r+1}(z))$ on level $r+1$ for sending a control output signal.

12. An apparatus according to Claim 11 wherein:
height $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_{r-1}, \dots, z_2, z_1, z_0]$ is converted to $h_r(z)$ on the level r by reversing the order of low-order z bits from z_{r-1} to z_0 into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2, \dots, z_{r-1}]$;
adding 1 (modulus 2^x); and
reversing back the low-order z bits; and
height z is converted to $H_r(z)$ on the level r by reversing the order of low-order z bits from z_{r-1} to z_0 into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2, \dots, z_{r-1}]$;
subtracting 1 (modulus 2^x); and
reversing back the low-order z bits.

13. An apparatus according to Claim 11 wherein:
height $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_{r-1}, \dots, z_2, z_1, z_0]$ is
converted to $h_r(z)$ on the level r by
reversing the order of low-order z bits from z_{r-1} to z_0
into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2,$
 $\dots, z_{r-1}]$;
adding J (modulus 2^r) in which J is an odd integer; and
reversing back the low-order z bits; and
height z is converted to $H_r(z)$ on the level r by
reversing the order of low-order z bits from z_{r-1} to z_0
into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2,$
 $\dots, z_{r-1}]$;
subtracting J (modulus 2^r); and
reversing back the low-order z bits.

14. An apparatus according to Claim 10 wherein a node
 $A(J, \hat{e}, z)$ on an outermost level J includes:

- a first interconnection with a device outside of the
multiple level structure for receiving message
data; and
- a second interconnection with a device outside of the
multiple level structure for sending a control
output signal.

15. An apparatus according to Claim 10 wherein a node
 $A(0, \hat{e}, z)$ on an innermost level 0 includes:

- a first interconnection with a device outside of the
multiple level structure for sending message data;
and
- a second interconnection with a device outside of the
multiple level structure for receiving a control
output signal.

16. An apparatus according to Claim 10 wherein:
on a level T , one ring is spanned in 2^T passes through
the angles \hat{e} from 0 to $K-1$ so that 2^T heights z
designate one ring.

17. An apparatus according to Claim 1, further comprising:
a plurality of devices coupled to the nodes of a level.

18. An apparatus according to Claim 1, further comprising:
a plurality of devices coupled to the nodes of level 0; and
a plurality of interconnect lines coupling the plurality of devices to respective nodes in the level J.

19. An apparatus according to Claim 18, wherein a device is coupled to a plurality of nodes in the level J.

20. An apparatus according to Claim 1, wherein:
 W_T rings are interconnected on a level T;
 W_{T-1} rings are interconnected on a level T-1; and
the W_{T-1} rings on level T-1 are divided into W_T mutually exclusive collections (C_1, C_2, \dots, C_{W_T}) such that each of the rings in collection C_n of level T-1 receive messages from ring R_m of level T.

21. A method of transmitting a message from a node N to a target destination in a first, a second and a third dimension of three dimensions in an interconnect structure arranged as a plurality of nodes in a topology of the three dimensions, the method comprising the steps of:

determining whether a node en route to the target destination in the first and second dimensions and advancing one level toward the destination level of the third dimension is blocked by another message;
advancing the message one level toward the destination level of the third dimension when the en route node is not blocked; and
moving the message in the first and second dimensions along a constant level in the third dimension otherwise.

22. A method according to Claim 21, further comprising the steps of:

specifying the first dimension to describe a plurality of levels, the second dimension to describe a plurality of nodes spanning a cross-section of a level, and the third dimension to describe a plurality of nodes in the cross-section of a level; sending a control signal from the node en route to the node N on a level q in the first dimension, the control signal specifying whether the node en route is blocked;

timing transmission of a message using a global clock specifying timing intervals to keep integral time modulus the number of nodes in a cross-section of a level, the global clock time interval being equal to the second time interval and the first time interval being smaller than the global time interval;

setting a first time interval \hat{a} for moving the message in the second and third dimensions;

setting a second time interval $\hat{a} - \hat{a}$ for advancing the message one level toward the destination level; and

setting a third time interval for sending the control signal from the node en route to the node N, the third time interval being equal to \hat{a} .

23. A method according to Claim 22, further comprising the steps of:

timing the message moving and advancing steps so that the messages enter node N on level q at times having the form $n\hat{a} + q\hat{a}$; and

timing the control signal sending step so that the control signals enter node N on level q at times having the form $n\hat{a} + q\hat{a}$ so long as the node en route is not blocked.

24. A method according to Claim 21, further comprising the steps of:

timing transmission of a message using a global clock;
setting a first time interval for moving the message in the second and third dimensions; and
setting a second time interval for advancing the message one level toward the destination level.

25. A method according to Claim 24, further comprising the steps of:

specifying the first dimension to describe a plurality of levels, the second dimension to describe a plurality of nodes spanning a cross-section of a level, and the third dimension to describe a plurality of nodes in the cross-section of a level;
specifying timing interval of the global clock to keep integral time modulus the number of nodes in a cross-section of a level, the global clock time interval being equal to the second time interval and the first time interval being smaller than the global time interval.

26. A method according to Claim 21 further comprising the steps of:

defining a header and a payload in the message;
encoding the destination in the second dimension in the header;
determining whether a potentially en route node is en route to the target destination including the steps of:
comparing the encoded destination in the second dimension to an encoded position of the potentially en route node;
resolving that the potentially en route node is en route when the encoded destination is the same as the encoded position of the potentially en route node.

27. A method according to Claim 26 wherein:
the destination in the third dimension in the header is encoded in a plurality of single-bit codes, each single-bit code relating to a level of the third dimension;
the position of the potentially en route node is encoded in a single-bit code; and
the comparing step is a single-bit comparison of the level-specific, single-bit destination code and the single-bit position code.

28. A method according to Claim 27 further comprising the step of:
discarding the level-specific, single-bit destination code in the as the message advances one level.

29. A method according to Claim 21 further comprising the step of:
on a level T , one ring is spanned in 2^T passes through the nodes in the second dimension so that 2^T nodes in the third dimension designate one ring.
interconnecting the three dimensional interconnect structure so that advancing of levels from a start level to the destination level furnishes access to all nodes in a ring.

30. A method according to Claim 21 wherein a message injected into the interconnect structure at a node $N(J, \hat{e}_1, z_1)$ and targeted to exit the interconnect structure at a node $N(0, \hat{e}_2, z_2)$ and injected at a time $(\hat{e}_2 - \hat{e}_1) \bmod K * \hat{a} + J\hat{a}$ causes the message to arrive at node $N(0, \hat{e}_2, z_2)$ at time 0.

31. A communication interconnect structure for transmitting messages, comprising:
a plurality of nodes arranged in a structure including:
a hierarchy of levels from a source level to a destination level;

- a plurality of nodes spanning a cross-section of a level; and
- a plurality of nodes in a cross-section span;
- a plurality of interconnect lines coupling the nodes in the structure including for a node N on a level L:
- a message input interconnect line coupled to a node on a previous level L+1;
- a message input interconnect line coupled to a node on the level L;
- a message output interconnect line coupled to a node on a subsequent level L-1; and
- a message output interconnect line coupled to a node on a subsequent level L-1.

32. An interconnect structure according to Claim 31, further comprising:

- a control input interconnect line coupled to the node on the subsequent level L-1 which is coupled to the message output interconnect line; and
- means for receiving a message on the control input interconnect line and, in accordance with the message, selectively transmitting a message on the message output interconnect line coupled to the subsequent level L-1 node or on the message output interconnect line coupled to the level L.

33. An interconnect structure according to Claim 32, further comprising:

- a control output interconnect line coupled to the node on the previous level L+1 which is coupled to the message input interconnect line;
- means for determining that a message is blocking the node N; and
- means for communicating via the control input interconnect line informing whether the node N is blocked.

34. An interconnect structure according to Claim 33, further comprising:

means for timing a message transmission time of a message transmitted from a level to a subsequent level and for timing a control signal transmission time of a control signal from a subsequent level to a level so that the control signal arrives first at a node.

35. An interconnect structure according to Claim 34, further comprising:

a control output interconnect line coupled to the node on the previous level L+1 which is coupled to the message input interconnect line;

means for determining that a message is blocking the node N; and

means for communicating via the control input interconnect line informing whether the node N is blocked.

36. A method of communicating messages in an interconnect structure comprising the steps of:

arranging a plurality of nodes in a structure including a plurality of hierarchical levels from a source level to a destination level, a plurality of nodes spanning a cross-section of a level and a plurality of nodes in a cross-section span, the nodes having an input connection on the same level, an input connection on a previous level, an output connection on the same level and an output connection on a subsequent level;

specifying a destination node in the destination level for receiving a message;

originating the message at a node in the source level;

communicating a message from node to node including the steps of:

determining at a node whether a node on a subsequent level is directed toward the destination node;

determining at a node whether the node on the subsequent level is blocked by another message;
advancing the message to the node on the subsequent level when the node is directed toward the destination node and a node is unblocked; and
otherwise traversing the message to a node on the same level.

37. A method according to Claim 36 wherein the step of determining whether a node on a subsequent level is directed toward the destination node further comprises the steps of:
encoding the destination node in a message in the header field;
encoding a designation of node position for the nodes at each level; and
determining that the node on the subsequent level is directed toward the destination node when the destination node encoding and the node position designation encoding correspond.